

SYSTEMS AND METHODS FOR INTERFERENCE MITIGATION AMONG MULTIPLE WLAN PROTOCOLS

This application claims priority to U.S. Provisional Application No.
5 60/238,761 filed October 6, 2000, the entirety of which is incorporated herein by
reference.

BACKGROUND OF THE INVENTION

The present invention is directed to systems and methods for preventing the
10 collisions or interference between signals from different wireless local area network
(WLAN) and wireless personal area network (WPAN) communication protocols that
coexist in the same frequency band. The term WLAN is used to refer to a class of
wireless communication technology that operates at a distance up to 100 meters, and
WPAN is commonly used to refer to a class of wireless communication technology that
15 operates up to a distance of 10 meters. For simplicity, when used herein, the term WLAN
is meant to encompass WLAN as well as WPAN technologies, and any other shorter-
range wireless communication technology, particularly, but not limited to, those that do
not require a license for operation by the Federal Communications Commission in the
United States and other similar unlicensed bands outside of the U.S.

20 The existence and popularity of new WLAN communication protocols results in
several protocols sharing the same frequency spectrum. This causes an interference
problem affecting throughput and reliability in wireless networks.

	IEEE 802.11b	Bluetooth™	HomeRF
Market Focus	Enterprise WLAN, school, home	Wireless cable	Home WLAN
Devices Likely to Use this Tech.	Laptop, desktop PCs	Palmtops, cell- phones, MP3 players	Home desktop PCs, printers, cordless phones
Technology	2.4 GHz ISM, DSSS	2.4 GHz ISM, FH	2.4 GHz ISM, FH
Peak Data Rate	11 Mbps	721 kbps	1.6 Mbps
Range	50m	10m	50m
Transmit Power	20 dbm	0 dbm	20 dbm

In the U.S. alone, for example, three popular WLAN technologies exist which share the 2.4 GHz unlicensed ISM band (see table above). As it turns out, each technology has merits relative to the other and, as a result, has secured itself as a preferred technology in at least one important market segment. Bluetooth™, for example, has
5 proven to be the leading WLAN technology for low-cost mobile computing devices such as palm-top computers, cell-phones, and MP3 audio players. On the other hand, because of its support for 10 Mbps data rates, IEEE 802.11b appears to be preferred for laptop computers in the enterprise and home environments.

The lack of a dominant WLAN technology means that different device types
10 may use multiple WLAN technologies at the same time in the same place. For example, in the residence, IEEE 802.11b and/or HomeRF may be used for wireless computer networking, and Bluetooth may be preferred for games, MP3 players and palmtop computers. In the enterprise environment, 802.11b may be used for wireless computer networking (laptops and/or desktop computers), and Bluetooth may be
15 used for palmtop computers and cell-phones.

Unfortunately, although each of these technologies was designed to work reliably with some interference in the ISM band, they were not designed to coexist with each other. The resulting interference between the WLAN technologies degrades system throughput and compromises the overall reliability of each of the
20 wireless networks. This is a well-known and well-documented problem.

SUMMARY OF THE INVENTION

The systems and methods of the present invention provide interference mitigation algorithms which allow for the operation of multiple wireless
25 communication protocols in a common frequency band, particularly an unlicensed frequency band allocated for short-range wireless communication. The systems and methods according to the present invention have utility in WLANs where there is possible overlap in frequency and time of signals transmitted in a common frequency band. In some cases, there is a signal of one communication protocol that is of a
30 fixed frequency nature (at least one or more fixed frequencies) in the frequency band concurrent with a signal of another communication protocol that hops to different, but predictable, frequencies. It is also possible that there are signals of a frequency hopping nature of the same type of communication protocol present in the frequency

band, for example, two Bluetooth networks operating in the same frequency band. In general, the system and methods of the present invention are useful in WLANs that have one or more frequency hopping communication protocols coexisting with one or more fixed frequency communication protocols in a frequency band, and WLANs that have two or more frequency hopping signals of the same or different type coexisting in a frequency band. Collisions of signals in the frequency band that would otherwise naturally occur among the devices are minimized or avoided while optimizing throughput of information.

Other objects and advantages of the present invention will become more readily apparent when reference is made to the following description in conjunction with the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram of a multi-protocol communication environment according to the present invention.

FIG. 2 is a block diagram showing an internal architecture suitable for a multi-protocol wireless communication device (MPD) according to the present invention.

FIG. 3 is a diagram depicting communication between an MPD and terminal nodes in a WLAN using two different types of WLAN communication protocols or technologies, referred to as A and B.

FIG. 4 is a frequency spectrum diagram for a multi-mode WLAN in which two different WLAN communication protocols operate in the same frequency band.

FIGS. 5A and 5B are timing diagrams showing basic time-division duplex formats for each WLAN communication protocol.

FIG. 6 is a diagram showing concurrent transmission of information from an MPD (downlink) on each of two WLAN communication protocols to two separate terminals in the same frequency band.

FIG. 7 is a diagram showing concurrent reception of information at the MPD (uplink) from two separate terminals on each of two WLAN communication protocols in the same frequency band.

FIG. 8 is a diagram showing transmission from the MPD to a terminal (downlink) on one communication protocol concurrent with reception at the MPD

from another terminal (uplink) on another communication protocol in the same frequency band.

FIG. 9 is a flow diagram showing a notch filtering technique to support simultaneous downlink of information on each of two different communication protocols in the same frequency band.

FIG. 10 is a flow diagram showing a notch filtering technique to support simultaneous downlink of information on one communication protocol and uplink on another communication protocol in the same frequency band.

FIG. 11A is a block diagram showing a configuration in an MPD suitable for simultaneous downlink of information using the frequency hopping communication protocol, such as Bluetooth™ and fixed frequency communication protocol, such as IEEE 802.11 in the same frequency band.

FIGs. 11B, 11C and 11D are spectrum diagrams of various signal points in the block diagram of FIG. 11A.

FIG. 12A is a block diagram showing a configuration in an MPD suitable for simultaneous downlink of information using a frequency hopping communication protocol and uplink of information using the fixed frequency communication protocol in the same frequency band.

FIGs. 12B and 12C are spectrum diagrams of various signal points in the block diagram of FIG. 12A.

FIG. 13 is a timing diagram showing the use of a guard packet to manage communication of synchronous type data using a frequency hopping protocol in the same frequency band as a fixed frequency communication protocol.

FIGs. 14A and 14B are flow diagrams showing a procedure for using the guard packet shown in FIG. 13 for downlink and uplink, respectively, of synchronous type data.

FIG. 15 is timing diagram for an alternative guard packet for wireless communication devices that function as access points.

FIG. 16 is a timing diagram for a guard packet for wireless communication devices that function as terminals or stations.

FIGs. 17A and 17B depict a flow diagram for a procedure for downlink communication of non-synchronous type data using a frequency hopping

communication protocol in the same frequency band as a fixed frequency communication protocol.

FIG. 18 is a spectrum diagram that shows an optional additional procedure for downlink communication of asynchronous type data using a frequency hopping communication protocol in the same frequency band as a fixed frequency communication protocol.

FIGs. 19A and 19B depict a flow diagram for a procedure for uplink communication of asynchronous type data using a frequency hopping communication protocol in the same frequency band as a fixed frequency communication protocol.

FIG. 20 is a timing diagram for a time frame of information transmitted in accordance with still another type of frequency hopping wireless communication protocol, such as HomeRF™.

FIG. 21A is a block diagram of a clock circuit arrangement designed to frequency lock and phase align the clock signals of two frequency hopping communication protocols.

FIG. 21B is a timing diagram showing a desired phase alignment between the clock signals of two frequency hopping protocols.

FIG. 22 is a diagram showing how an arbitration table is built to manage communication in the same frequency band where two frequency hopping communication protocols are active according to the present invention.

FIG. 23 is a flow diagram for a procedure to manage communication in the same frequency band where two frequency hopping communication protocols and one fixed frequency communication protocol are active and synchronous type data is communicated using the two frequency hopping communication protocols.

FIGs. 24A, 24B and 24C depict a flow diagram for a procedure to manage communication in the same frequency band where two frequency hopping communication protocols and one fixed frequency communication protocol are active and asynchronous type data is downlink communicated.

FIGs. 25A and 25B depict a flow diagram for a procedure to manage communication in the same frequency band where two frequency hopping communication protocols and one fixed frequency communication protocol are active and asynchronous type data is uplink communicated.

DETAILED DESCRIPTION OF THE INVENTION

1. General Description of Interference Mitigation Procedures

5

The present invention enables an improvement in the operation of two or more dissimilar wireless local area network (WLAN) communication protocols or technologies that operate in the same frequency band. For purposes of understanding the present invention, several terms will be introduced and explained. The term “node” will be used throughout. A node refers to a wireless communication device that is a point on a network, such as a WLAN. A “hub” node is a node that distributes information to (as well as receives information from) other nodes and may also be connected to another network, such as a wired network. A “terminal” node is a node that communicates with a hub node or other non-hub nodes, and is not connected to another network.

A multi-protocol wireless communication device (MPD) is a device which acts as a node (hub or terminal) on two or more wireless local or personal area networks, simultaneously. An example of an MPD is a device that operates simultaneously as a hub with respect to communication protocols A and B in communicating with terminal nodes. Some terminal nodes that communicate with this MPD may use only protocol A and others may use only protocol B, while still others may use protocol A and/or B. Generally, the devices that operate using communication protocol A form one wireless communication network, and the devices that operate using communication protocol B form another wireless communication network. Another example of an MPD is a laptop computer augmented with the appropriate hardware and/or software to act as a terminal node with respect to protocol A and simultaneously as a hub node with respect to protocol B. Still another example of an MPD is a device that simultaneously operates as a terminal node for protocol A and a terminal node for protocol B. The interference avoidance algorithms described herein are useful in the MPD. Finally, it should be understood that the interference avoidance algorithms of the present invention may also be useful to manage communication of two or more networks operating with the same communication protocol in the same frequency band, particularly where a

characteristic of the protocol (frequency hopping) necessitates collision avoidance procedures to optimize throughput on each network.

The terms “downlink” and “uplink” are used in the foregoing description. Transmission from the hub node to a terminal node is referred to as a “downlink” and communication from a terminal node to a hub node is referred to as an “uplink.” Transmission directly between terminal nodes is called peer-to-peer communication.

A terminal node may have a particular name or identifier for different communication protocols. For example, a BT terminal node is called a slave and a 802.11 terminal node is called a station (STA). A BT hub node is called a master. An 802.11 hub node is called an access point (AP).

FIG. 1 shows an example of a system 10 according to the invention in which an MPD 12 communicates with terminal nodes 14, 16, 18, 20 and 22. Terminal nodes 14 and 22 are wireless communication devices that, for example, operate using a first communication protocol, such as 802.11. Terminal nodes 16 and 20 are wireless communication devices that, for example, operate using a second communication protocol, such as Bluetooth. Terminal node 18 is a wireless communication device that, for example, operates using a third communication protocol, such as HomeRF. The MPD 12 is shown to be connected to a wired network 30, which in turn may be connected to a router 32 and to the Internet 34.

FIG. 2 shows an example of an internal architecture for an MPD device useful according to the invention.

The MPD 12 comprises a receive channel section 100, a transmit channel section 150, MAC layer protocol processors 180, 182 and 184, a network access arbitration controller 190, a voice processor 192 and a system controller 194. The receive channel section 100 comprises an RF-to-IF downconverter 102 which is coupled to a receive antenna, a digitizer 104 coupled to the downconverter 102, a plurality of IF-to-baseband downconverters 110, 112 and 114 each associated with a particular communication protocol, and a plurality of protocol detectors 120, 122, 124 each associated with a particular communication protocol.

In the transmit channel section 150, there is an IF-to-RF upconverter 152, a digital-to-analog converter 154, an adder 156, a plurality of baseband-to-IF upconverters 160, 162 and 164 each associated with a particular communication protocol and a plurality of modulators each 170, 172 and 174 each associated with a

particular communication protocol. Thus, for each communication protocol, there is a detector in the receive channel section 100 that performs the decoding of the baseband digital signal according to the rules of that communication protocol, and in the transmit channel section 150, there is a modulator that performs the encoding of the baseband digital signal according to the rules of that communication protocol.

For each protocol, the MPD 12 has a MAC block 180, 182 and 184. A network access arbitration controller (NAAC) 190 is coupled to the MAC blocks 180, 182 and 184. A voice processor 192 is coupled to the NAAC 190 to process digital voice data. (A video processor or any other special application processor may also be included in the MPD 12 as would be appreciated by those in the art.) A system controller 194 is coupled to the voice processor 194. The system controller 194 provides high level control functionality. For example, the system controller 194 can measure activity on one or more communication protocols in the network by monitoring signal energy in a frequency band of a particular or several communication protocols. Moreover, the system controller 194 can generate signals that represent a measure of a level of a signal. Many of the elements shown in FIG. 2, including the MAC blocks 180-184, voice processor 192, NAAC 190 and system controller 194 are software processes that are executed by one or more digital processing chips, either of an application specific variety or a general processing variety. The interference mitigation/collision avoidance procedures described herein are executed by the NAAC 190. However, it should be understood that other elements shown in FIG. 2 may share responsibility or perform the processes described herein. Moreover, the logic to perform the interference avoidance algorithms described herein can alternatively be embodied by one or more controllers configured or programmed (with software instructions or firmware), such as one or more general purpose processors or application specific integrated circuits (ASICs). Alternatively, the logic to perform the interference avoidance algorithms described herein can be embodied as a software product stored on a processor readable memory containing instructions that, when executed by a processor, cause the processor to perform the steps of the various algorithms described herein.

The precise location or component that performs the processes is not material to the invention. In some cases, such as the notch filtering procedures, the interference mitigation procedures involve tapping signals at the output of one or

more other elements in the receive section 100 or transmit section 150. Depending on design considerations, some components or functions may be implemented with analog signal designs, such as the upconverters and downconverters. On the other hand, digital signal filtering techniques may be desirable, but not required, such as in the case of notch filtering procedures described hereinafter.

With reference to the configuration shown in FIG. 3 in which an MPD communicates with node A (N_A) and node B (N_B). N_A and N_B use WLAN technologies (wireless communication protocols) A and B, respectively. N_A is located a distance of R_A from the MPD; N_B is located a distance of R_B from the MPD. The MPD is shown in the FIG. 3 to be connected to a wired network. In this example, the MPD is acting as a hub node with respect to protocol A and a hub node with respect to protocol B in the same frequency band.

Without loss of generality, the spectrum of wireless communication protocol A is assumed to occupy a greater bandwidth than that used by wireless communication protocol B, as shown in FIG. 4. The transmit power of protocol signal A and protocol signal B may also be different. Protocol signals A and/or B may be frequency-hopped throughout their allocated frequency band, and occasionally overlap in frequency.

Protocols A and B are further assumed to use a Time-Division Duplex (TDD) format, which typically results in a half-duplex mode of operation for each between the MPD and its nodes. As shown in FIGs. 5A and 5B, the nominal packet duration for communication protocol A is P_A . The minimum time between transmissions is G_A . Similarly, the nominal packet duration for WLAN technology B is P_B , and the minimum time between transmissions is G_B .

Interference occurs when two or more communication protocol signals overlap in time and frequency. The term collision is often used to describe the case where two or more signals attempt to occupy a common medium at the same time. As shown in FIGs. 6, 7 and 8 there are 3 general types of interference problems. (1) Downlink A – Downlink B (FIG. 6) occurs when the MPD wishes to transmit to N_A and N_B at the same time and protocol signals A and B at least partially overlap in spectrum. (2) Uplink A – Uplink B (FIG. 7) occurs when the MPD wishes to receive signals from N_A and N_B , where protocol signals A and B at least partially overlap in spectrum. (3) Downlink A – Uplink B (FIG. 8) occurs when the MPD wishes to

transmit a signal to N_A while receiving a signal from N_B , where both transmit and receive signals at least partially overlap in spectrum. A fourth case, Downlink B – Uplink A, is similar to (3).

1.1 Downlink A – Downlink B

When the spectrum of transmitted protocol signals A and B are non-overlapping, the MPD can deliver downlink protocol signals A and B to N_A and N_B , respectively, with no interference. When the spectrum of protocol signal A overlaps with that of B, conventional transmission techniques in which protocol signals A and B are transmitted without regard for the other will result in interference to both systems. With reference to FIG. 9, one method of communications that maximizes aggregate downlink throughput is to apply a notch filter 200 to the wider bandwidth protocol signal A to first remove the spectral energy of protocol signal A that overlaps with protocol signal B. The filtered signal A is then summed by an adder 210 with protocol signal B and transmitted to N_A and N_B .

A distortion will be introduced to protocol signal A when it is passed through a notch filter. N_A ideally uses a similar, matched filter to recover protocol signal A. N_A may apply a notch filter if N_A determines that an interfering signal occupies a portion of the spectrum used by protocol signal A, or if N_A is directed to do so by the MPD. If N_A is unable to provide a similar notch filter to its receive signal, the degree of distortion that may be tolerated without affecting link throughput will be dependent on the relative bandwidths of protocol signals A and B, and the robustness of the modulation of protocol signal A and error correction parameters.

If the spectral width of protocol signal A is comparable to protocol signal B, such that reliable communications cannot be achieved while simultaneously transmitting protocol signal B and protocol signal A, where signal A has been notch filtered as described above, the MPD should transmit exclusively only protocol signal A or protocol signal B. The choice of which WLAN to transmit may be determined by establishing a priority for different WLAN networks or for the type of data or channel that is to be carried on the WLAN. For example, voice and video communications, which are often carried on synchronous channels, are typically more sensitive to added latency than data communications. Voice and video channels would therefore normally be given a higher priority than given for a data exchange.

1.2 Uplink A – Uplink B

When the spectrum of transmitted protocol signals A and B are non-
5 overlapping, the MPD will be able to receive signals A and B from N_A and N_B ,
respectively, with no interference. When the spectrum of protocol signal A overlaps
with that of B, the two signals will interfere with each other and may lead to
incorrect decoding at the MPD of both signals. An improved method of
communications is to have the MPD keep such a situation from occurring by
10 exploiting Medium Access Control (MAC) properties typically available with short
range wireless communications systems to effectively hold-off or delay N_A or N_B
from transmitting, as described below.

A master-slave configuration is often used in wireless networks to provide
explicit control of when a node may transmit. The master typically must provide
15 explicit instruction to a slave before it may transmit data. If the MAC for protocol
signal A or protocol signal B uses a master-slave configuration, and the MPD is the
master, and the MPD can delay N_A or N_B from transmitting its uplink signal until
after the collision period.

Carrier Sense Multiple Access (CSMA) is often used to arbitrate wireless
20 network access by having nodes first determine that the wireless network is inactive
before transmitting. If protocol signals A or B use CSMA, the MPD may transmit a
downlink signal prior to the period of temporal and spectral overlap to keep other
nodes in the network from transmitting. This technique is particularly useful in
protecting packets delivered over isochronous channels that do not typically allow
25 retransmission. The MPD preferably transmits useful data on the downlink to one or
more nodes operating over A, for example, while protecting an uplink transmission
of protocol signal B in this case. In the event that the MPD has no useful data to
transmit to N_A , the MPD creates a dummy packet that contains no useful data for its
nodes, but serves the purpose of keeping N_A from transmitting an uplink signal. For
30 the case shown in FIG. 10 in which protocol signal A is transmitted by the MPD to
keep node N_A from transmitting, a notch filter 220 is applied to the wideband signal
A before it is transmitted to reduce the interference at the receiver of the MPD. For a
TDD system, a portion of the MPD transmit signal will be coupled into the MPD

receive path. The MPD can remove this coupled downlink signal by using a bandpass filter 225 and an adder 230 to pass only the desired uplink signal.

1.3 Downlink A – Uplink B

When the spectrum of protocol signals A and B are non-overlapping, the MPD will be able to transmit signal A and receive signal B from N_A and N_B , respectively, with no interference. When the spectrum of protocol signal A overlaps with that of B, the two signals will cause interference at the MPD and possibly at N_A as well. A method which improves the overall communications throughput is to delay or hold-off transmitting protocol signal A or B. A method that achieves still higher communications throughput is to apply a frequency notch on protocol signal A that corresponds to the bandwidth of protocol signal B, similar to that shown in FIG. 10.

1.4 Additional Collision Avoidance Techniques

Because wireless links may at times have a relatively high Bit Error Rate (BER) due to interference or low signal power, explicit acknowledgment (ACK) of the correct decoding of a data packet is often used as a mechanism to efficiently trigger a retransmission. For communication protocols that use an ACK as part of the communications handshake protocol, it becomes necessary to determine if there is a potential collision on the data transfer and the ACK. As an example, consider a data transfer over protocol A that includes the use of ACK in its protocol. Also consider that A and/or B is frequency-hopped throughout similar bands as described above. To ensure that the data exchange between the MPD and N_A is completed without interference, the techniques described above must be applied for both the data transfer and the ACK.

A technique referred to as polling may be used in a WLAN to enable contention-free data transfer. A slave node may only transmit after it has been polled by the MPD, which is assumed to be the WLAN master. The MPD may therefore use the poll as a mechanism to avoid collisions with other WLAN technologies. The MPD may also use polling with a combination of the above techniques, such as the

use of a frequency notch when the MPD has two different overlapping WLAN signals to transmit.

A technique that can enable higher data throughput in a frequency-hopped WLAN is to extend the length of a transmission for a given hop, and enable data to be transmitted during what would otherwise be a guard time. A guard time of G_B was earlier described with reference to FIGs. 5A and 5B. An additional benefit of extending a hop duration is to avoid hopping to a frequency which may interfere with a signal of another communication protocol. For example, if a protocol permits extending the hop duration to 5 hops, and hop 3 would be in the same band as another protocol signal, then the selection of a 5 hop transmission enables interference-free communications of data equal to 5 hops in duration. The same argument may be applied for a protocol that uses an ACK as part of its protocol; an extended hop duration is desirable if it enables both the forward channel or transmit signal and ACK to avoid interfering with another protocol signal.

The likelihood of two different protocol signals, such as A and B described above, interfering with each other is dependent on many factors. An important factor is the nominal network activity. If neither protocol A nor B are frequently active then the likelihood of interference between them is relatively small, and collision avoidance techniques may not be required. Conversely, if protocols A and B are relatively busy, there is a much higher likelihood that there will be significant interference then without the use of interference mitigation techniques may result in dramatic throughput reductions. The use of a metric that monitors WLAN activity can be used to determine if collision avoidance techniques should be deployed and if so, which ones to use. An additional metric that reflects the likelihood of a successful packet reception by a specific terminal may be used by the MPD to help determine whether a packet to that terminal should be transmitted during a period of potential interference.

A common technique to coordinate WLAN timing among multiple nodes is for the WLAN to use fixed duration slots, with an integral number of slots per frequency hop. For a situation in which an MPD operates 2 or more protocols that have a fixed slot structure, the MPD can enable further throughput enhancing techniques by time-aligning the signals of the protocols. Time-alignment of different

protocol signals enables the MPD to efficiently determine potential collisions and to take action to mitigate such collisions.

A common technique deployed in WLAN protocols to enable random access of the WLAN medium is often referred to as Carrier Sense Multiple Access (CSMA) in which nodes first determine if the channel is unoccupied before transmitting. Nodes typically measure the power in the channel to determine whether the channel is already in use. If an MPD is operating 2 or more communication protocols that use CSMA to access the WLAN, and there are periods where 2 or more protocol signals that occupy the same band, the MPD may use the techniques described above to mitigate potential interference. The MPD may also collect WLAN traffic metrics to arbitrate WLAN access between different protocols to improve total WLAN throughput and to help minimize access delay time in which terminals are delayed from using the WLAN medium while it is being used by other nodes. For those cases in which multiple communication protocols operate from an MPD and concurrently offer periods of network access via CSMA, a method that enables these protocols to have comparable levels of success in accessing the network is to provide similar configurations of the CSMA parameters, such as packet duration, guard time and back-off time. An MPD that simultaneously supports channels that do not allow retransmission, such as for voice or video, over 2 or more communication protocol technologies can improve overall network throughput by adjusting the time at which packets are to be transmitted on both the uplink and downlink to the extent permitted in the nodes of the protocols, to minimize any overlap of such time-sensitive transmissions.

2. Interference Mitigation Algorithms As Applied To Specific WLAN Protocols

2.1 Operation For MPDs Supporting a Frequency Hopping Protocol, such as Bluetooth and a Fixed Frequency Protocol, such as 802.11b DS

This section describes how the interference mitigation algorithm operates in networks supporting both Bluetooth (BT) and 802.11b DS (DS = Direct Sequence Spread Spectrum, 1-11 Mbps) protocols in the 2.4 GHz ISM band. Bluetooth is an

example of a frequency hopping protocol and 802.11 is an example of a fixed frequency protocol, both existing in the same frequency band.

The algorithm described in this section assumes that at most one BT piconet (a TDMA connection between a master and up to 8 slaves) and at most one channel of 802.11b DS are managed simultaneously from the MPD. However, the techniques described herein can be used to support multiple 2 WLAN communication protocols of the same type, e.g., two or more Bluetooth networks and two or more 802.11 networks.

2.1.1 802.11 Operation

An IEEE 802.11 WLAN network uses a wireless Ethernet-like protocol to pass data among its nodes. There are different physical layer (PHY) protocols that support frequency-hopped (FH) and direct sequence (DS) spread spectrum modulation formats. The interference mitigation algorithm described in this chapter assumes that the DS format is always used (although it is rather straightforward to modify the algorithm to support FH format using similar techniques).

The DS system specified in 802.11b supports data rates of 1, 2, 5.5 and 11 Mbps. BPSK modulation is used to support the 1 and 5.5 Mbps data rates; QPSK is used to support 2 and 11 Mbps.

The master of the 802.11 network is called an access point (AP). The other nodes are referred to as stations (STAs). An AP may also take on the attributes of a STA by providing them access to a distribution system that connects multiple wired or wireless LANs.

The medium access layer (MAC) protocol uses Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) to arbitrate data transfers among the nodes. To transmit a packet on the 802.11 network, the sending node first listens for a carrier on the 802.11 channel to see if it is currently being used. If the sending node detects that the channel is idle for at least 1 DIFS interval (DIFS = Distributed Coordination Function Interframe Space), it transmits its data. Otherwise, if the channel is busy, the sending node generates a random number to be used as a backoff counter. The sending node periodically listens to the channel and decrements the backoff counter during polling intervals in which the channel is idle. After the backoff counter reaches zero, the sending node transmits its message.

After receiving and correctly decoding the message from the sending node, the receiving node waits one Short Interframe Space (SIFS) period from the end of the message before transmitting an ACK message. The sending node waits up to one ACK timeout period, after which it concludes that its message was lost and executes a random backoff interval procedure before attempting to transmit its message, as described above.

Further details about the IEEE 802.11 protocols are found in the published ANSI/IEEE specification documents, entitled "ANSI/IEEE Std 802.11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," 1999 Edition, and subsequent versions thereof, which is incorporated herein by reference.

2.1.1.1 802.11 Network Configuration

In data only applications, the 802.11 network is generally configured to use large 802.11 fragment sizes (up to 1500 bytes) to achieve maximum system throughput. In combined voice/data applications, an MPD may generate control parameters to adjust the maximum fragment size of information transmitted by communication devices on the 802.11 network to minimize the required duration of a guard packet, which in turn minimizes the impact on throughput caused by the frequent deployment of 802.11 guard packets when synchronous type data on the Bluetooth network is to be transmitted. A description of guard packets is provided hereinafter.

2.1.1.2 Downlink of 802.11 Information

When the MPD has 802.11 downlink data to send, it transmits the data immediately using the rules described in the 802.11 MAC protocol. With reference to FIGs. 11A - 11D, a configuration is shown in which there is a 802.11 modulator 250, a BT modulator 255, a notch filter 260, an offset upconverter 265, an adder 270 and an upconverter 275. If at any time during the transmission the MPD also needs to transmit a BT slot in the 802.11 band, it notches a 1 MHz segment out of the baseband 802.11 signal spectrum to accommodate the transmit BT signal, adds the BT signal at the frequency location of the notch, attenuates the BT signal by at least

10 dB (in power) relative to the 802.11 signal, and transmits the composite signal. If the BT transmission terminates before the end of the 802.11 transmission, the segment removed from the 802.11 signal is replaced by an allpass filter having the same group delay as the notch. This approach allows the MPD to transmit both BT and 802.11 signals simultaneously in the same frequency band with little or no performance impact.

Indeed, since the standard BT terminal includes a 1 MHz lowpass filter preceding its FM discriminator, the filter will remove energy outside of the notch filter passband from the 802.11 signal, thus ensuring reliable data recovery in the BT receiver. Also, although the notch filtering and the addition of the BT signal cause some intersymbol interference and distortion of the 802.11 signal, simulations using a practical implementation of the notch filter have shown that the signal-to-distortion ratio of the composite signal when sampled at mid-chip to be at least 10 dB, enabling reliable data recovery in the standard 802.11 receiver.

2.1.1.3 Uplink of 802.11 Information

The MPD usually acts like a typical 802.11 AP when receiving 802.11 uplink data. The only exception to this occurs when the MPD receives 802.11 uplink data while simultaneously transmitting a BT downlink packet in the same band (for example, when the MPD transmits a BT POLL packet in the 802.11 band while receiving an 802.11 ACK). The configuration shown in FIGs. 12A-12C comprises a BT modulator 280, an upconverter 285, a power amplifier 287, an adder 290, a downconverter 292, a notch filter 295 and an 802.11 detector 297. In order to receive the 802.11 data reliably, the MPD notches out the transmit BT signal from the received 802.11 signal using a 1 MHz notch filter. The notch is replaced by an allpass filter having the same group delay as the notch as soon as the BT transmission terminates.

2.1.2 Bluetooth Operation

Current versions of Bluetooth uses two-level FSK modulation at a data rate of 1 Mbps. The FSK carrier is frequency hopped 1600 times per second to minimize interference from other sources in the 2.4 GHz ISM band.

Each 625-microsecond dwell defines a BT timeslot. The master sends downlink data to the slaves in the even-numbered timeslots; uplink data is sent from the slaves to the master in the odd-numbered timeslots.

Each frequency-hopped BT channel is associated with a *piconet*. The master of the piconet provides the hop sequence and the hop clock through which the master and slaves communicate.

Two data links are defined in the BT protocol – the Synchronous Connection Oriented (SCO) and the Asynchronous Connectionless (ACL). The SCO link carries full-duplex isochronous data such as digitized speech and/or video. SCO packets are single-slot packets and come in pairs – one packet for the downlink, the other for the uplink. SCO packets are never retransmitted.

The ACL link carries asynchronous data (e.g., for file transfers). ACL packets are 1, 3 or 5 slots in duration. A contention-free protocol is used to arbitrate ACL data transfers with a simple ACK/NACK retransmission scheme. The master polls each of the slaves for uplink data by sending a POLL packet in the appropriate timeslot. The slave addressed by the POLL packet responds with a 1, 3, or 5-slot uplink ACL packet in the following timeslot. The master sends an ACK/NACK to the slave in the next timeslot in which it addresses that slave.

Further details on the Bluetooth communication protocol can be found in the published specification documents by the Bluetooth Special Interest Group (SIG), entitled “Specification of the Bluetooth System,” Version 1.1, February 22 2001.

2.1.2.1 Bluetooth SCO Operation-Downlink and Uplink

This section describes how the MPD exchanges data with BT terminals using the SCO link. Although the SCO channel supports rate 2/3 and rate 1/3 FEC, interleaving is not supported, leaving the voice data vulnerable to bursts of bit errors whenever the BT hop frequency falls inside the 802.11 band. Consequently, in order to guarantee an acceptable level of voice quality in systems supporting both voice and data users, the MPD assigns a higher delivery priority to SCO packets than to ACL or 802.11 data.

With reference to FIG. 13, to guarantee reliable delivery of BT SCO data packets, the MPD transmits a “guard packet” on the 802.11 network immediately preceding the transmission of a BT SCO packet which falls inside the 802.11 band. The guard packet is an 802.11 message that alerts other 802.11 capable wireless

communication devices in the frequency band not to transmit information, and thereby allows the MPD to effectively take control of the 802.11 link immediately preceding the transmission of the SCO packet in order to guarantee its reliable delivery.

5 As FIG. 13 shows, the guard packet 300 has a header 305 and a trailer 310. The header includes an all-zero data portion 307 (all-zero information bits spread using the 11-bit barker code) followed by a short duration identifier (DURID) message 309. The all-zero portion 307 applies energy to the 802.11 link for a period equal to the duration of one 802.11 fragment, plus an SIFS interval, plus the duration
10 of an ACK message. This duration ensures that all 802.11 terminals will have stopped transmitting and entered the backoff state by the end of the all-zeros transmission. The DURID message 309 is a small data packet containing duration/id bit settings used to reserve the 802.11 bus for precisely the duration of the protected SCO packet.

15 The guard packet trailer 310 provides additional protection from 802.11 terminal transmissions during the SCO packet, and may also be used to carry live 802.11 downlink data if the MPD has transmit data waiting in its queues at the appropriate time. If the protected SCO packet is a downlink packet, the notch technique shown FIGs. 11A-11D is used to allow the MPD to transmit both the
20 downlink SCO data and the 802.11 guard packet trailer in the 802.11 band at the same time.

A similar technique is used if the SCO packet resides in the uplink. In this case, the MPD notches a 1 MHz segment from the transmitted 802.11 signal in the guard packet trailer to accommodate the BT signal being received. This prevents the
25 transmit signal energy from interfering with the BT signal at the BT receiver/detector input.

The implicit quality-of-service guarantee given to BT SCO voice channels over data channels comes at the expense of lower data throughput on the 802.11 network. The MPD network management software allows the network manager to
30 limit the number of allowed voice channels to upper-bound the throughput impact. A provision is also made to dynamically adjust the maximum number of allowed voice channels in proportion to the data network loading measured by the MPD.

The flow charts of the procedures for processing BT SCO data in the MPD are shown in FIGs. 14A and 14B, respectively. As the figures show, guard packets are only used when the SCO packets fall inside the 802.11 band. If at any time during the transmission of a downlink SCO packet the MPD also needs to transmit a 802.11 message, it notches a 1 MHz segment out of the 802.11 signal spectrum to accommodate the transmit BT signal, as described above. The MPD replaces the notch with an allpass filter having the same group delay as the notch filter when either the 802.11 or BT transmission terminates.

FIG. 15 illustrates an alternative form of a guard packet, useful in an MPD that functions as an 802.11 AP and as either a Bluetooth master or a Bluetooth slave. The header portion of guard packet 320 comprises a sequence of N 802.11 clear-to-send (CTS) packets 322. Each CTS packet 322, CTS(N) comprises a preamble 324, a frame control field 326, a duration ID (DURID) field 328, an address field 330 and a frame check sequence (FCS) field 332. The CTS packets are understood by 802.11 capable devices to not access the medium for a period of time. The DURID field 328 for any CTS packet CTS(i) is set to a value of $366 + (N-i)*L$, where L is the duration in μs of a CTS packet. For example, L is 304 μs using standard PHY preamble, and 1 Mbps transmission rate. The duration of a BT SCO slot is 366 μs . The number N of CTS packets in a guard packet is equal to $\text{ceil}(F/L) + 1$, where F is the fragment duration in μs at the lowest supported data rate, and L is defined above. Ceil (x) is a ceiling function meaning the smallest integer greater than (x). The fragment duration F is the longest transmit duration that a device can have on an 802.11 network. The theory or principle behind this alternative guard packet procedure is to guarantee that at least one CTS packet will be received by all of the 802.11 capable terminals by sending multiple packets over a time duration sufficient to extend longer than the longest transmission duration F that a device can have on the 802.11 network, regardless of the time alignment between the CTS packets and the other 802.11 transmission.

FIG. 16 illustrates yet another form of a guard packet, particularly useful for an MPD that functions as an 802.11 STA and either a Bluetooth master or a Bluetooth slave. The header portion of guard packet 340 comprises a sequence of subpackets (SPs) 342. Each SP 342 comprises an 802.11 request-to-send (RTS) packet 344 and a "no-TX" or silent interval 346. Each RTS packet 344 comprises a

PHY preamble 348, a frame control field 350, a DURID field 352 and an FCS field 356. An RTS packet is lower power than an CTS packet sent by an 802.11 AP. Like the procedure shown in FIG. 15, this procedure involves sending multiple RTS packets to guarantee that at least one RTS packet will be received by the 802.11 AP because several guard packets are sent for a time period that extends longer than the maximum 802.11 packet transmission permitted. When the 802.11 AP receives the RTS packet, it will, according to the 802.11 protocol, transmit a CTS packet in the frequency band with the proper parameters to reserve the medium for the 802.11 terminal node to receive or transmit the synchronous type data. The duration of the “no-TX” period is equal to the length of the SIFS plus the length of a CTS packet. The DURID field for RTS(i) packet is set to a value equal to $366 + (N-1)*L + M$, where L is the SP duration in μs , and M is the duration of a CTS packet plus an SIFS. The number N of SPs in a guard packet is equal to $\text{ceil}(F/L) + 1$, where F is the fragment duration as described above in connection with FIG. 15.

2.1.2.2 Bluetooth ACL Operation

This section describes how the MPD exchanges data with BT terminals using the ACL link. In the description that follows, it is assumed, without loss of generality, that the MPD has at least 5 packets of data in its queue to be transmitted to a BT terminal, allowing the MPD to use 5-slot packets to achieve the highest throughput efficiency. If less than 5 packets of data are available, single or triple-slot packets are transmitted using the approach described below.

2.1.2.2.1 Bluetooth ACL Operation-Downlink

The procedure for transmitting data on the downlink ACL from the MPD is shown in FIGs. 17A and 17B. First, if in step 400 the 802.11 network is idle and the network load is determined to be sufficiently small, the MPD in step 405 transmits a 5-slot ACL packet immediately, without regard to whether the BT hop frequency falls inside the 802.11 band at any time during the data exchange (i.e., during either the data transmission or the ACK). If the measured 802.11 network load is not sufficiently small, the MPD then in step 410 monitors the state of the BT hop frequency generator over the next several slots to see if it can complete a multi-slot transmission in which both the transmit data and the corresponding ACK are outside the 802.11 band. The MPD attempts to identify the largest multi-slot packet (1,3 or 5

slots) that satisfies this condition. If such a packet can be found, the MPD transmits the packet.

If such a packet cannot be found, then either the transmit frequency is in band, the ACK frequencies for the 1, 3 and 5 slot packets are all in-band, or both.

Therefore, steps 415 through 435 enumerated in FIGs. 17A and 17B perform the following logic. If the transmit frequency is in-band and at least one of the ACK frequencies is out-of-band, the MPD will transmit the shortest multi-slot packet that meets this condition if (1) the 802.11 network is idle for at least 1 DIFS, or (2) the MPD can guarantee that it will not receive 802.11 uplink data during the ACL packet's transmission, or (3) the MPD has been successful in transmitting T of the last ten ACL packets in the 802.11 band to the BT terminal during an active 802.11 uplink transmission, where T is an appropriate threshold. If none of conditions 1-3 can be satisfied, the MPD waits at least two slots before attempting to transmit again using the above procedure.

In case (3) above, the MPD prevents receive signal energy reflected from the transmitted BT signal from interfering with the 802.11 receiver by notching the BT signal from the receive signal path at the input to the 802.11 detector. The notch is replaced by an allpass filter having the same group delay as the notch as soon as the BT transmission terminates. This technique is illustrated in FIGs. 12A through 12C.

If at any time during the BT transmission the MPD also needs to transmit a 802.11 packet, it notches a 1 MHz segment out of the 802.11 signal spectrum to accommodate the BT signal, as described above in section 2.1.1.3. The MPD replaces the notch with an allpass filter having the same group delay as the notch when either the 802.11 or BT transmission terminates.

An optional and further refinement to the procedures of FIGs. 17A and 17B is shown in FIG. 18. The frequency band of the static protocol channel (e.g., an 802.11 channel) is shown at 454 and is centered at frequency $f_c(ST)$ Hz and extends between $f_c(ST) + X$ and $f_c(ST) - X$. Due to transmitter and receiver third order non-linear distortion, significant energy appears in frequency bands adjacent to the frequency band of the static protocol channel. This is called the IM3 band and is shown in FIG. 18 at 380 to exist between $f_c(ST) - 2X$ and $f_c(ST) - X$, and between $f_c(ST) + X$ and $f_c(ST) + 2X$. It is desirable to inhibit transmission of frequency hopping protocol signal 455 if its transmit frequency or its ACK at least partially

overlaps with the frequency band of the static protocol channel or with certain portions of the IM3 band caused by the non-linear distortion referred to above. The frequency of the forward/information packet using the frequency hopping protocol signal 455 is called $f(TX)$ and the frequency of the ACK packet is $f(RX)$.

5 The procedures of FIGs. 17A and 17B involve a determination whether $f(TX)$ and $f(RX)$ at least partially overlaps or falls within the static protocol channel between $f_c(ST) - X$ and $f_c(ST) + X$. Assuming the frequency(ies) of the forward channel packet(s) and the ACK are both outside the static protocol channel, the forward channel packet(s) will be sent (regardless of whether it overlaps with the IM3 band) if the $f(RX)$ is either outside the IM3 band completely or if it overlaps with the IM3 band, but the expected power of the ACK signal (determined from reception of prior ACK signals from that device) is greater than the level of the power in the IM3 band (at that frequency) by a required signal-to-noise ratio. If it is, then the forward channel packet(s) is transmitted using the frequency hopping protocol because it is known that the power of the ACK will be sufficient to be discriminated even in the presence of the IM3 energy. Otherwise, the packet is not transmitted, and transmission is deferred for two slots. That $f(TX)$ may fall within the IM3 band is not a concern because it is a signal that will be transmitted from the MPD at a sufficient power to be received by another device. In FIG. 18, frequency band 450 corresponds to the band where the frequency hopping protocol signal would be permitted to hop without using the optional procedures described herein. Reference numeral 452 identifies the expanded frequency band where the frequency hopping protocol signal is permitted to hop if the signal level of the ACK of the frequency hopping protocol signal is greater than the energy level of the IM3 band (by a predetermined amount) at that frequency (at the receiver). The wireless device performing this algorithm will have the capability of measuring the energy level of the IM3 band.

The use of the optional procedures of FIG. 18 together with the procedures of FIGs. 17A and 17B, achieves an overall logic flow that is summarized as follows.

Find largest n in $\{1,3,5\}$ such that:

1. The frequency of the forward channel packet(s), $f(TX)$, are outside of static protocol channel, and either

- 2(a) The frequency of the ACK signal, $f(RX)$ is outside of both the static protocol channel and the IM3 band, or
- 2(b) The frequency of the ACK signal, $f(RX)$ is inside the IM3 band and the expected power of the ACK signal exceeds the IM3 power (in the BT channel bandwidth) by an appropriate SNR threshold.

If such an n exists, then transmit; otherwise, wait two slots and then try again.

2.1.2.2.2 Bluetooth SCO Operation-Uplink

The procedure for receiving data on the uplink ACL at the MPD is shown in FIGs. 19A and 19B. First, in step 500, if the 802.11 network is idle and the network load is determined to be sufficiently small, in step 505 the MPD polls the BT slave without regard to whether the BT hop frequency falls inside the 802.11 band during either the POLL or uplink data packets. If the 802.11 network load is not sufficiently small, in steps 510 through 560 enumerated in FIGs. 19A and 19B, the logic performed is as follows. The MPD polls the slave only if (1) the hop frequency for the POLL packet and the uplink ACL packet are both outside the 802.11 band, or (2) the hop frequency for the POLL packet is inside the 802.11 band, the hop frequency for the uplink ACL packet is outside the 802.11 band, and either (a) the 802.11 network is idle for at least one DIFS interval at the beginning of the POLL packet's transmission, or (b) there is no 802.11 uplink data to receive at the MPD during the POLL packet's transmission, or (c) the MPD has been successful in transmitting T of the last ten POLL packets in the 802.11 band to the BT terminal during an active 802.11 uplink transmission, where T is an appropriate threshold. If conditions (1) and (2) are not met, the MPD waits at least 2 slots before attempting to use the above procedure again to receive uplink data.

2.2 Operation For MPDs Supporting Bluetooth, 802.11b DS and HomeRF

This section describes how the interference mitigation algorithm operates in networks supporting Bluetooth, 802.11b DS and HomeRF (up to 10 Mbps) protocols

in the 2.4 GHz ISM band. The approach is similar but slightly more complicated than the approach used to support only Bluetooth and 802.11 described above.

The algorithm described in this section assumes that at most one BT piconet (a TDMA connection between a master and up to 8 slaves), at most one channel of 802.11b DS, and one HomeRF network are managed simultaneously from the MPD. Again, the techniques described herein can be used to support multiple WLANs of the same type, e.g., two more Bluetooth networks, two or more 802.11 networks and two or more HomeRF networks.

2.2.1 HomeRF Operation

The HomeRF transmission protocol uses frequency-hopped FSK modulation (2 or 4 level FSK, up to 10 Mbps). The hop rate is 50 Hz. During each 20 ms dwell, the master (in this case, the MPD) and slaves exchange data in a superframe consisting of the following fields as shown in FIG. 20:

Beacon Interval – Period during which the master broadcasts superframe configuration data to the slaves (e.g., number and duration of each field, number of slots in each field, etc.)

CFP1 – Contention-free period used to retransmit lost voice packets from the CFP2 period of the previous dwell

Service Slot – Slot during which slaves are free to page the master to gain access to the network

Contention Period – Period during which the master and slaves exchange asynchronous data using CSMA/CA.

CFP2 – Contention-free period during which the master and slaves exchange isochronous voice information using TDMA.

Further details on the HomeRF communication protocol are found in the document published by the HomeRF Working Group, entitled “HomeRF 2.0 Protocol Specification,” 2000.

2.2.1.1 Timing

To prevent BT data transmissions from interfering with the HomeRF Beacon signal, the BT and HomeRF hop clocks are frequency-locked and phase-aligned in the MPD so that the Beacon Interval occurs during the “dead time” of the BT slot in

which there is no transmission, as shown in FIGs. 21A and 21B. Also (not shown in FIGs. 21A and 21B), the MPD generally attempts to phase align BT SCO transmissions with respect to the HomeRF superframe so that they do not overlap in time with HomeRF CFP1 voice slots (although this may not always be possible for various reasons). This minimizes the possibility of interruptions in synchronous data transmission, such as a voice or video, using both HomeRF and Bluetooth.

FIG. 21A shows a clock arrangement wherein a clock signal (CLK) from a common clock signal source is coupled to a counter 570. The counter 570 is in turn coupled to a comparator 572 and a comparator 574. The CLK signal is also coupled to a flip-flop 576. The output of the comparator 572 is coupled to the T input of the flip-flop 576 and to the reset input of the counter 570. The output of the comparator 574 is coupled to one input of an AND gate 578. The output of the flip-flop 576 is coupled to the other (inverted) input of the AND gate 578. The Q output of the flip-flop 576 is the Bluetooth clock signal (1600 Hz). The output of the AND gate 578 is coupled to the D input of a flip-flop 580. The Q output of the flip-flop 582 is coupled to the clock input of a modulo 16 counter 582. The TC output of the modulo 16 counter 582 is coupled to the T-input of flip-flop 584. The common CLK signal is coupled to the clock input of the flip-flop 584. The Q output of the flip-flop 588 corresponds to the HomeRF clock signal (50 Hz).

With reference to FIGs. 21A together with FIG. 21B, the operation of the clock arrangement is as follows. The CLK signal, which in the example, is a 2 MHz clock signal drives the counter 570 to count. The comparator 572 determines when the count value reaches 625, and when it does, it resets the counter 570, and drives the T input of the flip-flop 576. Thus, the output of the flip-flop 576 is 2 MHz divided by 625, a 1600 Hz clock signal useful for Bluetooth control. The comparator 574 determines when the count value of the counter 570 reaches 135 to drive one input of the AND gate 578. The output of the AND gate 578 will go high when the output of the flip-flop 576 is low concurrent with the output of the comparator 574 being high. This ensures that the AND gate 578 will go high only when Bluetooth clock signal is low. When the AND gate 578 goes high, the D input of the flip-flop 580 goes high, which in turn drives the count value of the modulo 16 counter 582 up one count, synchronized to the CLK signal. When the modulo 16 counter 582

reaches 15, it drives the T input of the flip-flop 584, which, synchronized to the CLK signal, outputs a clock signal useful for HomeRF control.

The result of the configuration shown in FIG. 21A is that the rising edge of the HomeRF clock is guaranteed to be 380 μ s (at least 366 μ s corresponding to the duration of a data or non-silent period of a Bluetooth time slot but not more than the period of the Bluetooth clock signal minus the duration of the Beacon interval) after the rising edge of the Bluetooth clock. Essentially, the logic generates the first and second clock signals by separately frequency dividing the common clock signal to generate first and second divided signals, and phase aligning rising edges of the first and second divided signals such that the rising edge of the first clock signal occurs a predetermined period of after the rising edge of the second clock signal. In this example, a convention is adopted in which the BT master schedules transmission of a time slot immediately following the rising edge of the Bluetooth clock and the HomeRF master schedules transmission of the Beacon interval immediately following the rising edge of the HomeRF clock. By ensuring that the rising edge of the HomeRF clock begins a time period equal to or greater than the duration of a Bluetooth slot, the Beacon interval, which starts a HomeRF superframe, will occur during a time period in which the Bluetooth network is quiet.

2.2.1.2 HomeRF Operation

The MPD adjusts its operating parameters before each 20 ms HomeRF dwell in order to minimize cross-network interference. A flowchart describing the MPD processing before each 20 ms HomeRF superframe is shown in FIG. 23.

If in step 600 the hop frequency for a particular HomeRF superframe falls outside the 802.11 band, the MPD configures the superframe normally (i.e., as it would in a HomeRF-only environment). For superframes in which the hop frequency falls inside the 802.11 band, the following steps are taken within block 610:

CFP1 and CFP2 are configured normally, except the voice slots are protected using an 802.11 guard packet (described above) to guarantee reliable voice data delivery.

Service slot configured normally.

Beacon signal configured normally.

Contention Period configured in one of two ways:

In one mode of operation - the so-called "shared CSMA mode", the 802.11 and HomeRF nodes are allowed to compete for access to the spectrum using CSMA/CA. In this case, the MPD must use similar network timing parameters for both networks (e.g., slot time, backoff interval, SIFS, DIFS, etc.) in order to ensure that both networks are given an equal opportunity to compete for access to the spectrum.

If Shared CSMA Mode is not enabled, the MPD alternates access ownership between 802.11 and HomeRF during the contention period based on their relative network load. For example, if the 802.11 network load is 10% and the HomeRF CSMA network load is 30%, the MPD allows the HomeRF to use the contention period 3 out of every 4 times both networks share the same frequency, and allows the 802.11 network to use the contention period 1 out of every 4 times. For superframes in which the MPD assigns ownership priority to HomeRF, the MPD transmits an 802.11 guard packet immediately before the contention period to prevent unwanted 802.11 network accesses. For superframes in which 802.11 is given ownership priority, the MPD removes the contention period from the superframe configuration data (in the beacon) to prevent HomeRF CSMA accesses.

To protect HomeRF voice slots from BT data collisions (in addition to using the phase alignment approach described above), the MPD uses a HomeRF/Bluetooth network arbitration table (arbTbl) to inhibit BT data transmissions that would otherwise interfere with the HomeRF voice slots. In general, the concepts of the arbTbl can be used to arbitrate the transmission of information of different data types (synchronous and asynchronous) and/or subtypes (voice, video, computer data) on two or more networks that may collide in frequency and time, where a priority is given to a certain data type on one network over the same or other data types on the other network. The arbTbl is shown in FIG. 22. There are 32 elements in the arbTbl. Each element corresponds to one of the 32 BT slots that comprise a 20 ms HomeRF superframe. If the MPD determines that a BT timeslot could collide with a HomeRF

voice data field, it sets the appropriate arbTbl element to one to inhibit transmission during that BT timeslot.

Before each new HomeRF hop to a new frequency, the MPD monitors the state of the BT hop frequency generator to determine the 32 BT hop frequencies for the next 20 ms HomeRF superframe. In block 620, for each BT hop frequency that interferes with the HomeRF band, the following action is taken:

If the BT slot overlaps with the CFP2 field: If the BT slot is a SCO slot, the MPD inhibits HomeRF transmission during the BT SCO slot by setting values in the Beacon interval to decommission (or eliminate) the appropriate CFP2 slot or slots for that superframe in the Beacon field. If the BT slot is not a SCO slot, the MPD inhibits the BT slot's transmission by setting the arbTbl element for that BT slot to one.

If the BT slot overlaps with the CFP1 field: If the BT slot is not a SCO slot, the MPD sets the appropriate arbTbl element to one in order to inhibit the BT slot's transmission. If the BT slot is a SCO slot, there are two procedures that are possible. One procedure, shown block 620 of FIG. 23, is to alternate priority between the BT SCO slot and the HomeRF CFP1 field. This is done by setting a variable, such as the one called "alternate" in the flow chart, to an initial value, for example, 1. Each time the logic proceeds through this process, the value of the alternate variable is toggled from 1 to 0, or 0 to 1. When the value of alternate is 1, the BT SCO transmission is disabled by writing a 1 to the arbTbl entry for that BT slot. When the value of alternate is 0, the values of the Beacon interval for the HomeRF superframe are set to decommission or eliminate the CFP1 field. This process of alternating priority between the HomeRF CFP1 and the BT SCO slot that collides with it, occurs with respect to similar collisions that occurred in prior superframes and similar collisions that may occur in future superframes. That is, if a similar collision was determined in a prior superframe and the BT SCO slot was given priority, then in the same type of collision in a current superframe would result in the HomeRF CFP1 having priority.

The second way of handling the situation when the BT slot is a SCO slot and overlaps with the CFP1 field (not specifically shown in FIG. 23) is to

attempt to place a non-information bearing "dummy segment" or segments at an appropriate location(s) in the superframe by manipulating the information in the Beacon interval for that the superframe in order to move the CFP1 field(s) away from the SCO slot. Alternatively, the values for the information in the Beacon interval can be set to ensure that the CFP1 field(s) do not coincide with the one or more BT time slots that are scheduled to transmit synchronous data. If for any reason it is not possible to move the CFP1 field away from those time slots, the MPD assigns alternating priority to the BT slot and the CFP1 field(s) in the manner described above using the alternate variable.

If the BT slot overlaps with the Contention Period: If the BT slot is a SCO slot, the MPD transmits a HomeRF guard packet (similar to the 802.11 guard packet described above) immediately before the SCO slot in order to block HomeRF CSMA transmissions during that slot.

The arbTb1 is built in advance for next 32 BT slots that coincide with the next HomeRF superframe. That is, prior to the transmission of the HomeRF superframe, an arbTb1 is built for the corresponding 32 BT slots. The value of the arbTb1 is read and examined prior to transmission of each BT slot. In addition, the parameters for a Beacon interval of a HomeRF superframe are determined prior to transmission of that superframe, so as to manipulate the position or existence of the CFP1 or CFP2 fields according to the logic of FIG. 23 and the foregoing description.

2.2.2 802.11 Operation

The approach used to manage 802.11 network activity from the MPD for networks supporting 802.11, HomeRF, and Bluetooth is identical to that used to support 802.11 and Bluetooth described above.

2.2.3 Bluetooth Operation

Bluetooth operation is similar, but slightly more complicated in multi-mode systems using 802.11, BT and HomeRF than in systems using 802.11 and Bluetooth only. These differences are discussed below.

2.2.3.1 SCO Operation

SCO operation is nearly identical to the 802.11/BT approach, except when HomeRF is involved, there is a chance that a SCO transmission could be disabled to

prevent a collision with a HomeRF voice slot. If the arbTbl entry for a downlink SCO slot is nonzero, the MPD does not transmit that slot. If the arbTbl entry for an uplink SCO slot is nonzero, the MPD directs the BT slave not to transmit the uplink SCO slot in the BT slot immediately preceding the SCO slot.

2.2.3.2 ACL Operation-Downlink

The procedure for transmitting data on the downlink ACL from the MPD is shown in FIGs. 24A, 24B and 24C. The approach is conceptually similar to the BT/802.11 approach; the only changes are the inclusion of (1) the arbTbl to inhibit ACL transmissions that would otherwise interfere with HomeRF voice slots, and (2) the HomeRF contention period in the CSMA processing.

First, in step 700, if the network loading of the 802.11 and HomeRF CSMA networks is determined to be sufficiently small, in step 705 the MPD identifies the largest multi-slot packet (1,3, or 5 slots) having non-zero arbTbl values during both the transmit portion and the ACK. If such a packet can be found, the MPD transmits the packet immediately in step 715, without regard to whether the BT hop frequency falls inside the 802.11 or HomeRF band during either the data transmission or the ACK. If such a packet cannot be found, in step 720 the MPD waits at least 2 slots before attempting to transmit ACL data again.

If the measured network load is not sufficiently small, in step 725 the MPD then monitors the state of the BT hop frequency generator over the next several slots to see if it can complete a multi-slot transmission in which (1) both the transmit data and the corresponding ACK are outside both active bands, and (2) none of the slots involved in the transmission (during either the transmit data or the ACK) have a non-zero arbTbl value. The MPD attempts to identify the largest multi-slot packet (1,3 or 5 slots) that satisfies this condition. If such a packet can be found, the MPD transmits the packet in step 730.

If such a packet cannot be found, then the last resort is to see if a packet duration can be found for which there are no nonzero arbTbl values during the transmit data or the ACK, the transmit frequency is inside either active band, and the ACK frequency is outside all active bands. The logic of steps 735 through 775 enumerated in FIGs. 24B and 24C is as follows. The MPD will transmit the shortest packet that meets these criteria if (1) the CSMA networks for each active band in which the ACL slot resides are idle for at least one DIFS interval, or (2) 802.11 is the

only active band in which the transmit slot resides and there is no 802.11 data to receive during the BT transmission, or (3) 802.11 is the only active band in which the transmit slot resides, there is 802.11 data to receive during the BT transmission, and the MPD has been successful in transmitting at least T of the last ten ACL packets in the 802.11 band to the BT terminal during an active 802.11 uplink transmission, where T is an appropriate threshold. If none of conditions 1-3 can be satisfied, the MPD waits at least 2 slots before attempting to transmit again using the above procedure.

In case (3) above, the MPD prevents receive signal energy reflected from the transmitted BT signal from interfering with the 802.11 receiver by notching the BT signal from the receive signal path at the input to the 802.11 detector. The notch is replaced by an allpass filter having the same group delay as the notch as soon as the BT transmission terminates. This technique is illustrated in FIGs. 12A through 12C.

If at any time during the BT transmission the MPD also needs to transmit a 802.11 packet, it notches a 1 MHz segment out of the 802.11 signal spectrum to accommodate the BT signal. The MPD replaces the notch with an allpass filter having the same group delay as the notch when either the 802.11 or BT transmission terminates.

2.2.3.3 ACL Operation-Uplink

The procedure for receiving data on the uplink ACL at the MPD is shown in FIGs. 25A and 25B. First, in step 800, if the arbTb1 values for the downlink POLL packet slot and the uplink ACL slot(s) are nonzero, the MPD does not send a POLL packet and waits at least 2 slots before attempting to receive uplink data in step 805.

Otherwise, if the network loading of the 802.11 and HomeRF CSMA networks is determined to be sufficiently small in step 810, in step 815 the MPD polls the BT slave without regard to whether the BT hop frequency falls inside the 802.11 or HomeRF band during either the POLL packet or the uplink ACL slot. If the network load is not sufficiently small, the logic of steps 820 through 870 enumerated in FIGs. 25A and 25B is performed as follows. The MPD polls the slave only if (1) the hop frequency for the POLL packet and the uplink ACL packet are both outside both active bands, or (2) the hop frequency for the POLL packet is inside either active band, the hop frequency for the uplink ACL packet is outside both active bands, and either (a) the active CSMA network(s) is idle for at least one

DIFS interval at the beginning of the POLL packet's transmission, or (b) if 802.11 is the only active network and there is no 802.11 uplink data to receive at the MPD during the POLL packet's transmission, or (c) if 802.11 is the only active network, there is 802.11 uplink data to receive during the POLL packet's transmission, and the
5 MPD has been successful in transmitting at least T of the last ten POLL packets in the 802.11 band to the BT terminal during an active 802.11 uplink transmission, where T is an appropriate threshold. If conditions (1) and (2) are not met, the MPD waits at least 2 slots before attempting to use the above procedure again to receive uplink data.

- 10 The interference avoidance procedures described above are useful individually or in combination. Moreover, each procedure or method may be embodied as a controller (one or more general or special purpose processors) configured or programmed to perform the various steps thereof, or as software stored on a processor readable memory containing instructions that, when executed by a
15 processor, causes the processor to perform the various steps thereof.

The above description is intended by way of example only.